

UNITED STATES PATENT APPLICATION

FOR

A METHOD AND APPARATUS FOR PROVIDING PROGRAMMABLE  
MEMORY FUNCTIONS FOR BI-DIRECTIONAL TRAFFIC IN A SWITCH  
PLATFORM

INVENTORS:

GENE CHUI  
LAMBERT FONG  
EUGENE WANG

PREPARED BY:

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP  
12400 WILSHIRE BOULEVARD  
SEVENTH FLOOR  
LOS ANGELES, CA 90025-1026

(408) 720-8598

Attorney's Docket No. 81862.P096

"Express Mail" mailing label number: EL034145605US

Date of Deposit: June 3, 1998

I hereby certify that I am causing this paper or fee to be deposited with the United States Postal Service  
"Express Mail Post Office to Addressee" service on the date indicated above and that this paper or fee has  
been addressed to the Assistant Commissioner for Patents, Washington, D. C. 20231

Cindy Baglietto

(Typed or printed name of person mailing paper or fee)

(Signature of person mailing paper or fee)

(Date signed)

jc518 U.S. PTO

09/090096



12400 WILSHIRE BOULEVARD

A METHOD AND APPARATUS FOR PROVIDING PROGRAMMABLE  
MEMORY FUNCTIONS FOR BI-DIRECTIONAL TRAFFIC IN A SWITCH  
PLATFORM

5    FIELD OF THE INVENTION

The present invention relates generally to the management of network nodes and, more particularly, to a mechanism for managing cell traffic in a multiservice switch platform.

10   BACKGROUND OF THE INVENTION

Until recently there has persisted a fundamental dichotomy between different types of telecommunication networks. A first type of telecommunication network, the telephone network, switches and transports predominantly voice, facsimile, and modulation-demodulation system (modem) traffic. A second type of telecommunication network, the data network, switches or routes and transports data between computers.

Telephone networks were developed and deployed earlier, followed by data networks. Telephone network infrastructures are ubiquitous, however, and as a result data networks typically are built, to a limited extent, using some components of telephone networks. For example, the end user access link to a data network in some cases is implemented with a dial-up telephone line. The dial-up telephone line thus connects the end user computer equipment to the data network access gateway. Also, high speed digital trunks interconnecting remote switches and routers of a data network are often leased from telephone carriers.

Nonetheless, telephone and data network infrastructures are usually deployed together with limited sharing of resources, especially with regards to the core components of the networks--the switches and routers that steer the

payloads throughout the networks. The cost of this redundancy coupled with advances in data network technologies has led, where possible, to integrated data traffic comprising voice, data, facsimile, and modem information over a unified data network. As such, a data network should now be able to accept,  
5 service, and deliver any type of data on a random, dynamic basis using a minimum set of hardware on a single platform.

Multiservice network switches are used to provide a data path, or interface, between multiple networks, each of which may operate using a different type of data or according to a different networking standard protocol.

10 Examples of the networking protocols supported by these multiservice switches include, but are not limited to, frame relay, voice, circuit emulation, T1 channelized, T3 channelized, and Asynchronous Transfer Mode (ATM).

Typical prior art switch platforms handling data of different types, speeds, and bandwidths exhibit a problem in that there is a significant risk of  
15 data loss because of the different data parameters. In order to reduce the risk of data loss, the typical prior art switches use first-in-first-out (FIFO) buffers. The problems with the typical prior art FIFOs are numerous but, overall, can be generalized as inflexible and performance limiting. First, the typical prior art FIFOs are limiting in that they are unidirectional and only support data  
20 flow in one direction. Next, the interface through the typical prior art FIFO is synchronous, thereby limiting the types of data accepted by the switch platform and the types of subscriber equipment coupled to the platform. Furthermore, the typical prior art FIFO has a fixed-width data path.

The typical prior art FIFO exhibits congestion problems in that it fails to  
25 provide adequate information regarding the quantity and contents of the FIFO. For example, the prior art FIFO does not provide quantitative information regarding the availability of cell space in the FIFO, and the cell

space availability indications typically provided are based on word boundaries. Additionally, the prior art FIFO fails to provide a way to discard invalid data that has been written to a FIFO. Furthermore, the typical prior art FIFO limits diagnosis and switch platform performance monitoring as it fails to provide for a non-destructive processor read of the FIFO or a diagnostic loopback between multiple FIFOs.

In addition to the aforementioned problems with typical prior art FIFOs, these FIFOs fail to provide fair rate-based bandwidth allocation among cell traffic sources. Instead, the typical prior art FIFOs use a round-robin priority scheme to service the traffic sources, wherein all traffic sources are serviced in consecutive order with equal bandwidth being allocated to each traffic source. In a wide area network, for example, bandwidth allocation is important because customers are serviced and charged according to the bandwidth used. Therefore, a bandwidth allocation scheme that provides equal bandwidth to all traffic sources regardless of source requirements is inefficient and cost ineffective.

Furthermore, the typical prior art FIFO is problematic in that it fails to provide adequate cell traffic congestion management, a failure that results in dropped and lost cells. Furthermore, the prior art FIFO inefficiently manages the routing of information through a switch platform. This inefficiency results from globally assigning the cell bus logical connection number (LCN) and the UDF the same value for each connection. This global assignment tightly couples the LCN address space and the UDF address space, thereby forcing a large portion of the cell bus address space to be unusable.

Furthermore, this global assignment requires switch platform software to make queries to multiple switch cards to find an unallocated value for use in both the LCN and the UDF address spaces.

## SUMMARY AND OBJECTS OF THE INVENTION

It is therefore an object of the invention to provide fair rate-based cell traffic arbitration and bandwidth allocation between multiple cell traffic sources.

5 It is a further object of the invention to prevent dropped or lost cells in a switch platform by providing cell bus traffic congestion management.

It is a further object of the invention to provide flexibility and a performance improvement in the translation of cell traffic routing information.

10 These and other objects of the invention are provided by a parameterized bi-directional FIFO unit that controls cell traffic in a switch platform using a first and a second unidirectional FIFO buffer. The first and second unidirectional FIFO buffers each comprise asynchronous read and write ports. A cell size and a word size of the first and second unidirectional  
15 FIFO buffers are programmable. The bi-directional FIFO unit is coupled to write at least one cell from and read at least one cell to at least one asynchronous transfer mode (ATM) interface, at least one frame relay interface, at least one voice interface, and at least one data interface. As such, the first unidirectional FIFO buffer is coupled to write at least one cell from an  
20 ATM interface, a frame relay interface, a voice interface, and a data interface. The first unidirectional FIFO buffer is coupled to read at least one cell to at least one switch, wherein the switch handles cells from sources having a number of bandwidths. The switch is coupled to route the at least one cell to an OC12 trunk line and to at least one service module. The service module is  
25 coupled to provide the cell to at least one service subscriber using T1, E1, T3, E3, OC3, and OC 12 ports.

The second unidirectional FIFO buffer is coupled to read at least one cell to an ATM interface, a frame relay interface, a voice interface, and a data interface. Furthermore, the second unidirectional FIFO buffer is coupled to write at least one cell from at least one switch, wherein the switch handles

5 cells from sources having a number of bandwidths. The switch may be coupled to route the cell from an OC12 trunk line and from a service module. The service module is coupled to provide the cell to at least one service subscriber using T1, E1, T3, E3, OC3, and OC 12 ports. Invalid cells are discarded from each unidirectional FIFO buffer.

10 Other objects, features, and advantages of the invention will be apparent from the accompanying drawings and from the detailed description which follows below.

## BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is illustrated by way of example and not limitation in the figures of the accompanying drawings, in which like references indicate similar elements, and in which:

5        **Figure 1** is a multi-service switch platform of one embodiment. The switch platform comprises two Processor/Switch Modules (PSMs).

**Figure 2** is a switch platform of one embodiment as a stand alone node.

**Figure 3** shows switch nodes of one embodiment connected as a network through router interfaces.

10       **Figure 4** shows switch nodes of one embodiment connected as a network through ATM interfaces.

**Figure 5** is a switch of one embodiment as a node in a BPX/IPX network.

15       **Figure 6** is a system block diagram of the switch platform of one embodiment.

**Figure 7** is a PSM comprising Cell Bus Controllers (CBCs) of one embodiment.

**Figure 8** is a cell format of the QE mode of one embodiment.

**Figure 9** is a cell format of the ATM mode of one embodiment.

20       **Figure 10** is a block diagram of a Cell Bus Controller (CBC) of one embodiment.

**Figure 11** is a diagram of an egress data flow of a CBC of one embodiment.

25       **Figure 12** is a diagram of an ingress data flow of a CBC of one embodiment.

**Figure 13** is a block diagram of an Egress Cell Input Controller (ECIC) of one embodiment.

**Figure 14** is a block diagram of an Egress Cell Processor (ECP) of one embodiment.

**Figure 15** is a block diagram of a Multicast Engine (MCE) of one embodiment.

5      **Figure 16** is a block diagram of a Cell Bus Master (CBM) of one embodiment.

**Figure 17** is a block diagram of a Cell Bus Slave (CBS) of one embodiment.

**Figure 18** shows the CBM cell bus signal timing of one embodiment.

10      **Figure 19** shows the CBS cell bus signal timing of one embodiment.

**Figure 20** is a block diagram of an Ingress Cell Processor (ICP) of one embodiment.

**Figure 21** is a block diagram of an Ingress Cell Output Controller (ICOC) of one embodiment.

15      **Figure 22** is a flowchart of a method for routing at least one cell in a switch platform of one embodiment.

**Figure 23** is an address map RAM data field definition of a CBC of one embodiment.

20      **Figure 24** is a mapping for a first PSM card in a first slot of one embodiment.

**Figure 25** is a mapping for a second PSM card in a second slot of one embodiment.

**Figure 26** is a service module configuration of one embodiment in which service module A communicates with service module B.

25      **Figure 27** is a service module configuration of one embodiment in which a service module fails and a redundant service module is on the same CBC as the failed service module.



**Figure 28** is a service module configuration of one embodiment in which a service module fails and a redundant service module is on a different CBC from the failed service module.

**Figure 29** is a parameterized bi-directional synchronous cell FIFO with  
5 loopback and processor access interface implemented in the CBC of one embodiment.

**Figure 30** lists the configuration parameters of the bi-directional FIFOs of a CBC of one embodiment.

**Figure 31** lists the input/outputs (I/Os) and a corresponding interface  
10 description for the bi-directional FIFOs of a CBC of one embodiment.

**Figure 32** is a flowchart for controlling cell traffic in a switch platform of one embodiment.

**Figure 33** is a block diagram of the flag generation of one embodiment.

**Figure 34** is a timing diagram of a write port of one embodiment.

**Figure 35** is a timing diagram of a read port of one embodiment.  
15

**Figure 36** is a flowchart for providing cell bus management in a switch platform of one embodiment.

**Figure 37** is a timing diagram of a processor port of one embodiment.

**Figure 38** is a bandwidth distribution of one embodiment using a  
20 parallel resistor network model.

**Figure 39** is a flowchart for providing cell traffic arbitration in a switch platform of one embodiment.

**Figure 40** is a block diagram of an interface loopback of one embodiment.

**Figure 41** is a block diagram of an individual cell bus loopback of one  
25 embodiment.

## DETAILED DESCRIPTION

Intended advantages of the invention described herein may comprise the provision of synchronous parameterized FIFO memory functions for bi-directional data flow through a switch platform, asynchronous read and write ports in a unidirectional FIFO, a FIFO having a programmable cell size, and a FIFO having a programmable word size, but the embodiment is not so limited. Furthermore, intended advantages may comprise the provision of a FIFO that provides quantitative information based on cell boundaries regarding the availability of cell space in the FIFO, a way to discard invalid data written to a FIFO, a non-destructive processor read of a FIFO, and a diagnostic loopback between multiple FIFOs, but the embodiment is not so limited.

**Figure 1** is a multi-service switch platform 100 of one embodiment. The switch platform 100 comprises two Processor/Switch Modules (PSMs) 102-104. Each PSM 102-104 is coupled to communicate between a backplane 110 and eight service modules (SMs) 120-154, but the embodiment is not so limited. The outside interfaces provided by the service modules 120-154 and the PSMs 102-104 provide the user a large selection of interfaces at speeds from DS0 modem access to OC3 ATM interface comprising T1/E1, DS3, OC3, modem, and router communications, but the embodiment is not so limited. The PSM cards 102-104 provide high speed native ATM interfaces through the backplane 110, wherein the interfaces comprise T3, OC3, and OC12 interfaces which can be configured as ATM ports or trunks, but the embodiment is not so limited. In one embodiment, the PSMs 102-104 support switched connections between service module ports and between PSM back card ports. The PSMs 102-104 are coupled to the service modules 120-154 using a service module interface with the physical layer convergence function

performed on the PSM. The PSM embodiments support 1.2 Gbps, 15 Gbps, and 45 Gbps thereby allowing support for each of these rates on the same switch node, but the embodiments are not so limited. The higher capacity PSM embodiments may be deployed in configurations where there are  
5 broadband interfaces in the network, and trunks can be supplied from a broadband interface.

The switch platform 100 of one embodiment supports applications requiring a combination of switching, routing, and other interface functions. As such, the switch platform may stand alone or function as part of a larger  
10 network. Therefore, the switch platform described herein may be configured either as a routing node or a feeder node attached to a router node, but the embodiment is not so limited. As a routing node, the platform will support AutoRoute and PNNI simultaneously.

**Figure 2** is a switch platform 200 of one embodiment as a stand alone  
15 node. Furthermore, multiple nodes can be coupled to form a network. The trunking of the nodes can be through the ATM interfaces 202 provided on the service modules of each switch platform, the ATM interfaces provided on the PSM 204, or a wide variety of interfaces, for example, a 100mbps Ethernet connection provided through the router interface. **Figure 3** shows switch  
20 nodes 302-306 of one embodiment connected as a network through router interfaces 312-316. **Figure 4** shows switch nodes 402-406 of one embodiment connected as a network through ATM interfaces 412-416. Furthermore, the ATM interfaces on the switch nodes of one embodiment provide ATM native connections to other types of network switches. **Figure 5** is a switch platform  
25 502 of one embodiment as a node in a BPX/IPX network 504.

**Figure 6** is a system block diagram of the switch platform 600 of one embodiment. As previously discussed herein, the switch platform 600

comprises two PSMs 602-604 and sixteen service modules, wherein two service module front cards 606-608 each comprise four service modules and two service module back cards 610-612 each comprise four service modules, but the embodiment is not so limited. The PSMs 602-604 provide, but are not limited to, the following functions: switch fabric; processor; cell bus interface to service modules; clocking; ethernet and RS232 ports; and system environment monitoring. The PSMs 602-604 provide a shared memory switch fabric capable of sustaining two times OC12 bandwidth. Furthermore, the PSMs 602-604 provide the switch platform 600 control functions and a switch fabric for intra-nodal cell switching. Moreover, the PSMs 602-604 provide a variety of interfaces for satisfying trunk requirements. In one embodiment, the system bandwidth may be increased by changing the PSM, thereby allowing for the use of multiple high bandwidth service modules.

The service modules of one embodiment support services including, but not limited to, Internet Protocol (IP), frame relay (FR), circuit emulation, and ATM. The narrow band interfaces, or service modules, supported in one embodiment comprise, but are not limited to, eight port channelized and unchannelized T1/E1 frame relay service modules, eight port T1/E1 ATM service modules, and four port circuit emulation. Furthermore, supported service modules comprise eight port CO-FRAD, eight port T1/E1 circuit emulation, T3/E3 circuit emulation, two port 16M HSSI FR, two port DS1 and DS0 channelized DS3 FR, two port DS3 unchannelized DS3 FR, two port 52M HSSI FR, ATM service modules, and voice card service modules. Moreover, broadband interfaces, or service modules, are supported.

In one embodiment, each PSM supports 128k cell storage within the switch, but the embodiment is not so limited. In one embodiment, the total bandwidth on the switch is more than two OC12, wherein the bandwidth is

partitioned for local switching and for trunk switching. Local switching refers to switching between service modules of a switch platform while trunk switching refers to switching between a PSM trunk and either the service modules or the PSM trunk. Alternate embodiments may support an aggregate of 24 Gbps bandwidth.

Eight cell buses are supported by each PSM of one embodiment. Each cell bus is capable of sustaining one OC3 bi-directional bus; however, the PSM may support a two times clock mode providing more than OC6 capacity. The cell buses provide a high speed interface between the switch fabric and the service modules using a Cell Bus Controller ASIC (CBC ASIC or CBC). The CBC interfaces with the eight OC3 cell bus lanes provided by the PSM. To facilitate communication between PSMs, a cell bus slave is added to each PSM. In one embodiment, the cell bus slave on a first PSM of the switch platform is coupled to one of the cell buses driven by a second PSM of the switch platform. Therefore, a PSM may send a cell back to itself through the switch. In one embodiment, the local switching function supports 16K bi-directional connections, and the trunk switching function supports 32K bi-directional connections, but the embodiment is not so limited.

The cell bus of one embodiment comprises a number of buses that provide inter-connectivity and redundancy to the switch modules. In particular, the cell bus provides a high speed cell data path between the service modules and the PSM switch fabric. The cell bus comprises a collection of independently controlled bus lanes. Each of the bus lanes, or sub-cell buses, provides up to two OC3 bandwidth, wherein a total of four cell bus lanes are provided for each sub-shelf. The four cell bus lanes, where each is capable of supporting one OC3 bandwidth, allows the embodiment to

sustain two times OC12 bandwidth. Furthermore, two sets of cell bus lanes are provided for supporting service module redundancy.

The switch platform software of one embodiment comprises several software subsystems distributed across the PSM and the service modules. The software subsystems are separated into three layers: node management software; network control software; and platform control software.

Subsystems in each of these layers are separated by application programmer interfaces (APIs). A virtual switch interface (VSI) serves as a standardized interface on the switch thereby allowing one controller implementation to be used across different platforms for the same function. The switch platform of one embodiment supports VSI slaves for PNNI, AutoRoute, and tag switching. Within a switch platform, VSI-related functions are performed by one or more processors known as VSI slaves. The VSI slaves may be organized using a centralized model, a distributed model and a hybrid model; each of these models may be supported by the platform of one embodiment, but the embodiment is not so limited.

**Figure 7** is a PSM 700 comprising Cell Bus Controllers (CBCs) 702-704 of one embodiment. Each CBC 702-704 controls cell traffic between four Cell Bus Masters (CBMs) CBM0-CBM7, one Cell Bus Slave (CBS) CBS0-CBS1, and two Queue Engine ASICs (QE ASIC or QE) 712-714, but the embodiment is not so limited. It is noted that, as used herein, the term "ingress" is used to refer to cell traffic flow toward the switch 712-714, or QE, and the term "egress" is used to refer to cell traffic flow away from the switch 712-714. Thus, the cell traffic flow from the QE 712-714 to the CBC 702-704 is termed egress, and the cell traffic flow from the CBC 702-704 to the QE 712-714 is termed ingress. The CBC 702-704, in acting as a master to four cell buses and a slave to one cell bus, controls cell traffic between the five cell buses and the two QEs 712-714. A

CBC interface couples the CBCs 702-704 and the QEs 712-714. In one embodiment, a CBC can handle 32 devices on the CBC interface, but the embodiment is not so limited.

The CBC of one embodiment supports one OC-12 (622 Mbit/s) cell input/output (I/O) bandwidth, meaning that the CBC can read/write a cell from/to the QE within approximately 680 nanoseconds, but the embodiment is not so limited. Each CBM supports an OC-3 bandwidth for transmitting and receiving. The CBC of one embodiment may implement cell storage using first-in-first-out (FIFO) buffers along the cell datapath. The CBC is used twice in the QE mode in each PSM card. Each CBC receives cells over a bus from one QE but sends cells over the bus to both QEs. On egress, the CBC routes traffic to the CBS.

Each CBM of one embodiment comprises clock logic, which can operate at either 21 MHz, 42 MHz, or 50 MHz, and a clock input pin, but the embodiment is not so limited. Moreover, all four CBMs can run synchronously on one common clock. Each CBS of one embodiment handles bus traffic of speeds up to OC-3; therefore, the aggregate bandwidth for the cell buses is more than one OC-12. Each CBS of one embodiment runs on a 21 MHz, 42 MHz, or 50 MHz clock, but the embodiment is not so limited.

The CBC of one embodiment supports two modes, a QE mode and an ATM mode. In supporting the QE mode, the CBC utilizes two transmit and two receive slave ports of the CBC interface. The CBC interface uses a 16-bit data path that transmits cells that each comprise twenty-seven 16-bit words. **Figure 8** is a cell format of the QE mode of one embodiment. The two CBC interface slave ports, which have separate control signals and one common datapath, are used to send cells to the two QEs, where data is sent to only one QE at a time. The transmit slave port is used to receive cells from one QE.

In supporting the ATM mode, one CBC interface receive slave port and one CBC interface transmit slave port are used, but the embodiment is not so limited. The CBC interface comprises an 8-bit datapath that transmits cells that each comprise 56 byte words. **Figure 9** is a cell format of the ATM mode of one embodiment. Egress traffic is routed to the CBS in the ATM mode.

**Figure 10** is a block diagram of a Cell Bus Controller (CBC) 1000 of one embodiment. **Figure 11** is a diagram of an egress data flow of a CBC 1000 of one embodiment. **Figure 12** is a diagram of an ingress data flow of a CBC 1000 of one embodiment. The CBC 1000 comprises, but is not limited to, an Egress Cell Input Controller (ECIC) 1002, an Egress Cell Processor (ECP) 1004, a Multicast Engine (MCE) 1006, a CBM 1008, a CBS 1010, an Ingress Cell Processor (ICP) 1012, and an Ingress Cell Output Controller (ICOC) 1014.

**Figure 13** is a block diagram of an Egress Cell Input Controller (ECIC) of one embodiment. In operation the ECIC receives cells from one QE 1302, using the cell bus transmit slave interface. As such, the ECIC is the first module of one embodiment through which incoming cells pass on the egress datapath of the CBC. The primary interfaces of the ECIC are with the QE or ATM 1302 and the ECP 1304. The ECIC is responsible for interfacing with the QE/ATM 1302 over a cell bus interface, wherein the QE/ATM 1302 is the cell bus transmit master and the ECIC is the cell bus transmit slave. Furthermore, the ECIC stores ATM cells sent by the QE 1302 in a two-cell deep FIFO 1306, and provides ATM cell availability status to the ECP 1304. Moreover, the ECIC converts the 16-bit input datapath in the QE mode, or the 8-bit input datapath in the ATM mode, to a 32-bit internal datapath.

**Figure 14** is a block diagram of an Egress Cell Processor (ECP) of one embodiment. The ECP is the second module of one embodiment through which incoming cells pass on the egress datapath of the CBC. The primary



interfaces of the ECP are with the ECIC 1402, the MCE 1404, the CBS, four CBMs 1408, and the ICOC 1410. The ECP processes cells received from the QE via the ECIC 1402 and forwards the cells to one of four CBMs 1408, a CBS, or an MCE 1404. Processing includes Cell Bus Header look-up from the Cell Bus Header RAM 1412 for cell bus traffic and Multicast (MC) Address RAM 1414 lookup for multicast traffic. For ATM traffic, the ECP simply forwards the cell to the CBS, without the Cell Bus Header look-up. The ECP unloads cells from the ECIC input FIFOs and checks the parity of the data read from these FIFOs.

The ECP performs cell bus interface polling address to target address mapping using thirty-two possible polling address destinations that are mapped to twenty-six targets. The mapping is performed by Address Map RAM 1416 look-up logic, wherein the targets are the service modules via a CBM, a second PSM Card via a CBS, and a MCE.

For cell bus cells in the QE mode of one embodiment, the ECP generates a Cell Bus Header (CBH) random access memory (RAM) address for the cells using a logical connection number (LCN), accesses the CBH RAM for the CBH entry, and generates a new Logical Slot ID in the CBH, wherein the physical slot ID is found in the Address Map RAM. Furthermore, the ECP adds the CBH to the cells and sends the cells to one of the CBMs or the CBS, depending on the address. If the number of cells in the egress FIFO of the target reaches a threshold, the ECP discards the incoming cell. The threshold is defined as the maximum number of cells in the FIFO, 8 cells in one embodiment, with a 2-bit threshold configuration subtracted, but the embodiment is not so limited.

For MC cells in the QE mode of one embodiment, the ECP uses the upper twelve bits of the LCN to perform Multicast Address Look-up. The ECP replaces the first UDF in the cell with the MC Record RAM Address and sends

the cells to the MCE. If a back pressure signal from the MCE is active, the cell is discarded by the ECP.

In the ATM mode of one embodiment, cells are routed to the CBS without any processing, except for a parity check. When a loopback control is active, all traffic is routed to the ICOC. When the cell available signal from the ICOC is inactive, the cell is discarded. Any parity error in the cell header will cause the cell to be discarded, while any parity error in the payload will set an error flag but will not cause the cell to be discarded.

**Figure 15** is a block diagram of a Multicast Engine block (MCE) of one embodiment. The MCE and the associated queues comprise the third module through which MC cells pass on the egress datapath of the CBC. The primary interfaces of the MCE are with the ECP 1502, the ICP, and the ICOC 1504. The MCE is responsible for storing MC cells in either the critical 1506 or the non-critical queue 1508, replicating and sending MC cells for each LCN in a given MC Group at a rate of up to OC-12 to the ICOC 1504. The MCE sends cells to the ICOC 1504 when the ICP has no cells to send. The MCE provides queue storage for thirty-two cells, of which sixteen are critical cells and sixteen are non-critical cells. The MCE replicates MC cells for each LCN associated with that PSM card in a given MC Group at a rate of up to OC-12, and inserts a 5-bit target slot ID, read from a MC record RAM 1510, into the fourth word of the cell. The MCE sends replicated cells to the ICOC 1504 when datapath grants are given by the ICP, and generates back pressure to the ECP 1502 when either of the critical 1506 or non-critical queues 1508 reach a threshold that is configured by firmware.

**Figure 16** is a block diagram of a Cell Bus Master (CBM) 1600 of one embodiment. The CBM 1600 interfaces with the cell bus 1602 as a master; as such, the CBM 1600 handles cell traffic in both ingress and egress directions.

The CBM 1600 is used four times in the CBC of one embodiment, but the embodiment is not so limited.

The CBM 1600 and associated FIFOs 1610-1612 are used to handle both ingress and egress cell bus traffic. During data ingress, the CBM 1600 arbitrates requests from service modules on the cell bus 1602 and grants one service module the right to send a cell in the next cell time. In arbitrating, the CBM 1600 provides polling and bandwidth specific arbitration logic to determine which service module can send cells to the CBM 1600. The CBM grant logic grants CBM access to the service module selected so the service module can send a cell to the CBM 1600. In supporting data ingress, the CBM 1600 provides a dedicated eight-cell ingress FIFO 1612 for storing a maximum of eight cell bus cells, but the embodiment is not so limited. The CBM 1600 interfaces with the ICP via the CBM ingress FIFO 1612; the CBM 1600 writes to the ingress FIFO 1612 while the ICP reads from the FIFO 1612 and processes the read cell. The CBM 1600 checks parity over each 8-bit word as the word is read from the cell bus 1602, converts the received cell to a 32-bit word, and writes the 32-bit word into the ingress FIFO 1612. Moreover, the CBM 1600 checks the slot ID of the cell received against the slot ID of the service module that was granted access to the CBM 1600 in order to verify that the correct service module is sending the cell. The CBM 1600 polls for status comprising back pressure, ready, and present status, but the embodiment is not so limited.

During data egress, the CBM 1600 provides a dedicated 8-cell egress FIFO 1610 for storing a maximum of eight cell bus cells. The ECP writes to the FIFO 1610 while the CBM 1600 reads from the FIFO 1610. The CBM 1600 reads 32-bit words from the egress FIFO 1610 and checks parity over each of the two 16-bit words read. The CBM 1600 converts each 16-bit word to two 8-bit words and generates parity. Each 8-bit word and the corresponding parity is

provided to the cell bus 1602. As the read cell is transmitted to all service modules on the cell bus 1602, each service module will read in the cell and determine if the cell is destined for that service module by looking at the broadcast field and the slot ID field in the Cell Bus Header. The CBM 1600  
5 checks to make sure a service module acknowledges receipt of the transmitted cell. The CBM 1600 has the capability to reset each service module.

Figure 17 is a block diagram of a Cell Bus Slave (CBS) 1700 of one embodiment. The CBS 1700 interfaces with the cell bus 1702 as a slave, wherein the CBS 1700 and associated FIFOs 1710-1712 are used to handle both  
10 ingress and egress cell bus traffic. The CBS 1700 interfaces with the ECP and the ICOC. On ingress, the CBS 1700 receives cells from the cell bus 1702, wherein the CBS 1700 provides a dedicated 16-cell ingress FIFO 1712 for storing a maximum of 16 cell bus cells. The CBS 1700 writes to the FIFO 1712 while the ICP reads from the FIFO 1712. Furthermore, the CBS 1700 checks  
15 parity over each 8-bit word as the word is read from the cell bus 1702. The CBS 1700 checks the Slot ID and the MC bit in the Cell Bus Header of the cell received against the Slot ID of the CBS 1700 and accepts the cell if the Slot ID matches or the MC bit is asserted; otherwise, the cell is discarded. As a cell is received, the cell is converted to a 32-bit word before being written into the  
20 ingress FIFO 1712. An acknowledge signal is sent to the transmitting CBM after the cell is received. The accepted cell is written to the CBS ingress FIFO 1712. The ICP reads the cell from the ingress FIFO 1712 and processes the cell.

On egress, the CBS 1700 is responsible for providing a dedicated 8-cell egress FIFO 1710 for storing a maximum of eight cell bus cells. The ECP writes  
25 to the egress FIFO 1710 while the CBS 1700 reads from the egress FIFO 1710. The CBS 1700 converts a 32-bit word read from the egress FIFO 1710 to two 16-bit words, and checks parity over each 16-bit word. The CBS 1700 makes a

request to an associated CBM during a poll phase to send a cell to the cell bus 1702, and waits for a grant. If a grant is received, the CBS 1700 sends the cell to the cell bus 1702 and the CBM. Furthermore, the CBS 1700 generates ready, present, and back pressure status. Moreover, on egress the CBS 1700 provides the capability to loopback egress cell traffic, wherein the loopback provides cell traffic to ingress for diagnostic individual cell bus loopback. **Figure 18** shows the CBM cell bus signal timing of one embodiment. **Figure 19** shows the CBS cell bus signal timing of one embodiment.

**Figure 20** is a block diagram of an Ingress Cell Processor (ICP) of one embodiment. The ICP is the second module through which incoming cell bus cells pass on the ingress datapath of the CBC. The primary interfaces of the ICP comprise the four CBMs, the CBS, the ICOC 2004, and the MCE 2002. The ICP is responsible for determining service decisions by arbitrating between four CBMs and one CBS. Once the service decision is determined, a cell is then read from the FIFO of the selected source and processed. If the CBC is in the QE mode, the ICP will do a look-up in the external LCN RAM for the LCN and insert it into the cell. No processing is done in the ATM mode of one embodiment, but the embodiment is not so limited. An ACP, or processor, can manipulate the priority of CBS traffic by setting a 6-bit CBS priority configuration. The ICP unloads cell bus cells from the ingress cell bus FIFO of the selected source, and checks parity over the data read from the input FIFO.

In the QE mode, the ICP removes the Cell Bus Headers from the cell bus cells and forms a new cell. The ICP generates LCN RAM addresses, looks up LCNs in the external LCN RAM using the generated address, inserts the new LCN in the new cell, extracts the target slot ID from the Cell Bus Header

and inserts it into the new cell, and transmits the generated new cells to the ICOC.

In the ATM mode, the ICP transfers the cell to the ICOC after a parity check. When there is no cell to send in a cell time, the ICP grants the MCE  
5 the cell time so the MCE can send the cell to the ICOC.

**Figure 21** is a block diagram of an Ingress Cell Output Controller (ICOC) of one embodiment. The ICOC is the last module through which ATM cells pass on the ingress datapath of the CBC, but the embodiment is not so limited. The primary ICOC interfaces comprise the interfaces with the ICP  
10 module, the ECP module, the MCE module, and the QE 2102 or ATM. The ICOC determines the QE destination from the target slot ID and sends cells to that QE using the cell bus receive slave interface.

The ICOC stores ATM frames sent by the MCE, the ICP, or the ECP into the two-cell FIFO 2104. The ICOC interfaces with the ingress QE/ATM 2102  
15 over a cell bus interface. Furthermore, the ICOC converts the 32-bit internal datapath to a 16-bit interface datapath in the QE mode, and to an 8-bit interface path in the ATM mode.

An ACP interface and control logic is responsible for handling ACP requests to the CBC. The primary interfaces of the ACP interface and control  
20 logic are with the external ACP sub-system and the CBC internal modules. The ACP interface and control module is responsible for monitoring the ACP interface and looking for valid CBC requests from the ACP. The ACP interface and control module maintains various control and status registers, and provides configuration information to the CBC. Furthermore, the ACP  
25 interface and control module performs ACP-requested actions and passes the required information to the appropriate module so that it can complete the operation.

Cell bus logic controls cell bus cell flow to and from the five external cell buses of one embodiment. The cell bus logic is implemented in the CBM and the CBS. Each of the four CBMs of one embodiment control an associated cell bus. The CBS can interface with one of the four cell buses controlled by the four CBMs or it can communicate with a dedicated cell bus. In the PSM, the CBS communicates over a dedicated cell bus. Furthermore, the CBS is connected to the cell bus which is controlled by a CBM in another PSM card; this is how the two PSMs communicate with each other.

When the CBM grants a slave module request to send a cell to the requesting slave module during the next cell time, it presents the grant address and asserts a grant enable bit from cell bus cycle 51 to cycle 58. When the slave receives its Physical Slot ID on the grant bus along with the asserted grant enable signal any time during the grant cycle, the slave is enabled to transmit a cell onto the bus during the next cell time.

In the operation of a slave module sending data to a CBM, when the slave module which has made a request to send a cell to the CBM receives the grant from the CBM, the slave module starts sending the cell at the beginning of the next cell bus frame time following assertion of the Tx-Frame signal at cycle 1. The first byte of the cell is sent to the CBM at cycle 1, while the last byte of the cell is sent at cycle 56.

As the CBM is receiving data from the slave module, the CBM checks to see if the cell is coming from the slave module to which the grant was provided during the last cell bus frame time. If a mis-comparison occurred, an error condition is indicated. The CBM also checks for parity on each byte of the cell as it is received into the CBM.

In the operation of one embodiment of a CBM sending data to slave modules, the first byte of a cell is transmitted in two cycles, cycle 58 and cycle

1. When the CBM is ready to send a cell to the slave module, it will begin sending the cell starting at the last cycle of the current cell bus frame time, cycle 58. The first byte of the cell is sent out again at the first cycle of the next cell bus frame time, cycle 1. The second byte of the cell is sent out in cycle 2, while the last byte of the cell is sent out in cycle 56. In cycle 57, the CBM sends out all zero data. Parity generation is performed on all cell data bytes.

In one embodiment, transmitted cells are sent to all slave modules on the cell bus. Each slave module must read in the first byte of the cell and determine if the cell is addressed to the slave module. As such, the slave module determines whether the physical slot ID field in the first byte of the cell matches the physical slot ID of the slave module receiving the cell, or determines whether the MC bit in the first byte of the cell is a 1. If either of these conditions are satisfied the cell is read and processed; otherwise, the cell is discard.

The slave modules can send three types of status to the CBM, but the embodiment is not so limited. The three types of status are sent using the poll bits like the slave module requests, wherein the poll bits are asserted at different cycles for each different status. The three types of status comprise: a ready status, wherein the slave module is working and ready to interface with the CBM; a present status, wherein the slave module is present but may not be ready; and a stop status, wherein the slave module is congested and cannot accept any more cells from the CBM.

The CBC of one embodiment provides cell bus traffic congestion management between the CBMs and the CBSs of a switch platform. In performing cell traffic congestion management, a CBS determines that it is congested and is running out of buffer space to store the incoming cells. The congested CBS informs, or back-pressures, the corresponding CBM to stop



sending cells to the congested CBS. The back-pressure notification uses a back-pressure bit defined in the cell bus interface for each service module on a cell bus. When the back-pressure bit is asserted for a service module, the CBC will prevent the QE from sending cells to that service module. This notification should be sent before the CBS FIFO is full and completely runs out of space to store incoming cells so as to provide time for the CBM to react while providing enough storage space for the cells already in transit to the CBS. Upon receipt of the notification from the CBS, the CBM ceases transmission of cells to the congested CBS. The congested CBS will update the congestion notification to the CBM when the CBS ceases to be congested. The CBM, in turn, continuously monitors the congestion notification and resumes sending cells to the CBS when the CBS ceases to be congested.

In order to provide the traffic management function without adding more signals to the cell bus, the congestion notification is performed using an existing request and status poll signal. When the CBS is congested it will assert a poll signal at cell bus cycle 50, when the CBS is in an odd slot of the cell bus frame, or cycle 56, when the slave is in an even slot of the cell bus frame. In one embodiment, CBS congestion is indicated when the CBS input FIFO has space for less than eight cells, wherein the eight cells handle CBM reaction latency and cells already in transit, but the embodiment is not so limited. Upon receipt of the congestion notification by the CBM, the CBM forwards the congestion notification message to the ECP. As a result, the ECIC blocks cell transmission to the CBC ASIC.

The CBC of one embodiment provides flexibility and performance improvement in the translation of routing information between source and target service modules. In the QE mode the CBC supports CBH look-up in the ingress direction, or cell bus to QE direction, and the egress direction, or QE to

cell bus direction. The CBH RAM contains one primary and one alternate CBH for each of the 64K logical connection numbers (LCNs). The logical organization of the RAM is 128K x 17, but the embodiment is not so limited.

In operation, cells are routed among service modules by the CBC using  
5 a first memory to convert an address having a first format into an address having a second format. In one embodiment, the address having the first format is received in a header of a cell, and the address format comprises a 17-bit cell bus logical connection number of a destination port. The address having the second format is a 16-bit UDF used by a switch of the switch  
10 platform. The address having the first format is used to form a third address that is used to access the first memory. The data located at the third address of the first memory is a 16-bit UDF used to address the switch.

A second memory is used to convert an address having the second  
15 format into an address having the first format. The address having the second format is used as a fourth address to access the second memory. The data located at the fourth address of the second memory is a 32-bit cell bus header that addresses a destination port. The destination port is at least one service module of the switch platform. The first and second memory comprise an external random access memory, but the embodiment is not so  
20 limited.

In one embodiment, each connection of a service module has a 17-bit cell bus LCN comprising a 12-bit LCN, a 4-bit Slot ID, and a cell bus number. As the QE, or the switch, has a 16-bit UDF, a mapping is used to map between the LCN and the UDF. The mapping is used because when a cell is  
25 transmitted to the CBM on ingress to be switched by the QE, the 17-bit cell bus LCN is converted to the 16-bit UDF used by the QE. Following switching, when a cell is enroute to the target service module on egress, the 16-bit UDF is

converted to the 17-bit cell bus LCN for the target service module. The conversions in both the ingress and the egress direction are performed by the CBC via an external RAM lookup. The CBC uses two external memory structures, the CBH RAM and the LCN RAM.

5           On cell ingress, the 12-bit cell bus LCN and the 4-bit Slot ID from the cell bus Header of the incoming Cell Bus cell are extracted. Furthermore, the cell bus number, indicating whether the cell is coming from CBM 0, 1, 2, 3, or the CBS, is formed by mapping a 3-bit cell bus source number to a 1-bit cell bus number. This mapping may be programmed by firmware. The 12-bit LCN, 4-  
10 bit Slot ID, and the cell bus number are used to form a 17-bit address to access the external LCN RAM. The data from the LCN RAM is the 16-bit UDF and it is forward along with the cell to the QE.

          On cell egress, the 16-bit UDF of the incoming QE cell is used as the address for the external CBH RAM. The data of the CBH RAM is a 32-bit  
15 CBH. It is added to the cell and sent to the service module.

**Figure 22** is a flowchart of a method for routing at least one cell in a switch platform of one embodiment. Operation begins at step 2202, at which a logical connection number and a slot identification are extracted from a first cell header of a cell received from at least one service port. A cell bus number  
20 is formed, at step 2204, by mapping a cell bus source number of the first cell header. A first address is formed, at step 2206, from the logical connection number, the slot identification, and the cell bus number. The first address may comprise a 17-bit address, but the embodiment is not so limited. A first external memory is accessed using the first address, at step 2208. A 16-bit UDF  
25 is read from the first external memory, at step 2210, and the UDF is forwarded to the switch, at step 2212.

Furthermore, a cell may be routed from a switch to a service module, or port, by addressing a second external memory, or CBH RAM, using a 16-bit UDF of at least one cell received from a switch. A second cell header is read from the second external memory, and the second cell header is added to the cell. The second cell header may comprise 32 bits, but the embodiment is not so limited. The second cell along with the second cell header is sent to a destination service port.

The CBC of one embodiment provides 1-to-1 and 1-to-N service module redundancy. Therefore, when any service module in the system fails, a redundant module can be brought on-line to replace the failed module. This redundancy is provided using two CBCs having a total of 128K of Cell Bus Headers and supporting 64K connections. For each connection, there is a primary CBH entry in one CBC, and an alternate CBH entry in the other CBC. The contents are the same but the addresses of these two CBHs are different by one bit, the alternate bit. Each service module has an alternate bit. When a new connection is made, the firmware adds a primary CBH entry to the CBC that currently has the target service module, and adds an alternate CBH entry to the other CBC; this is done to make sure CBH information is on both CBCs. The firmware chooses which of the two CBH entries to use by setting the alternate bit in the Address Map RAM.

**Figure 23** is an address map RAM data field definition of the CBC of one embodiment. Up to two PSM cards are supported in two slots of a switch, but the embodiment is not so limited. A set of mappings is defined and, based on the mappings, the firmware programs an Address Map RAM and the Cell Bus Master Configuration Registers. **Figure 24** is a mapping for a first PSM card in a first slot of one embodiment. **Figure 25** is a mapping for a second PSM card in a second slot of one embodiment.

**Figure 26** is a service module configuration of one embodiment in which service module A (SM A) communicates with service module B (SM B). In this example, SM A and SM B are controlled by the same CBC CBC1, but each SM could be controlled by different CBCs. When one service module sends cell traffic, or traffic, to another service module, the traffic goes through a series of steps in being routed. The primary CBH entry is used. The cell traffic is received by the CBC controlling the source service module. In the CBH, the 5-bit Target Slot ID specifies the target service module. The CBC looks into the QE destination register to find out to which QE the traffic should be routed; the traffic is then routed to that QE. Once the current QE receives the cell traffic, it will forward the cell traffic to the proper CBC, depending on its LCN entry; the CBC uses an address to index into the Address Map RAM for this information. The CBC uses the LCN number to get the CBH from the CBH RAM. When a service module fails, the redundant service module can be either on the same CBC or on another CBC.

**Figure 27** is a service module configuration of one embodiment in which a service module SM B fails and the redundant service module D (SM D) is on the same CBC CBC1 as the failed service module SM B. When SM B fails, the redundant SM D takes over, wherein SM D is controlled by the same CBC CBC1 as SM B. In replacing SM B with SM D, SM D is initialized with the information resident in SM B, and the Address Map RAM is changed so that the address for SM B is now pointing to SM D.

**Figure 28** is a service module configuration of one embodiment in which a service module SM B fails and the redundant service module C (SM C) is on a different CBC CBC2 from the failed service module. When SM B fails, the redundant SM C takes over, wherein SM C is controlled by a different CBC CBC2 than SM B. In replacing SM B with SM C, SM C is

initialized with the information resident in SM B. In CBC1, the QE-destination register is modified to have traffic to SM B forwarded to QE2. In CBC2, the Address Map RAM is modified to have traffic intended for SM B forwarded to SM C. In CBC2, the Address Map RAM is modified to use the  
5 alternate CBH entry. In CBC2, the QE-destination register is modified to have traffic to SM A forwarded to QE1.

In one embodiment, each CBC supports 26 devices comprising 24 cell bus service modules (8 service modules per cell bus), a PSM Card, and a Multicast Engine, but the embodiment is not so limited. As the  
10 corresponding firmware and software deals with the CBC devices in a 5-bit Logic Slot ID, a mapping of Logical Slot ID to Cell Bus Number and Physical Slot ID is used, but the embodiment is not so limited. Furthermore, the firmware and software transmits cells to the CBC ASIC via the QE. As the QE Chip has 16 polling addresses, another level of mapping is used.

15 A parameterized bi-directional synchronous cell FIFO with loopback and processor access interface is used for the ingress and egress FIFOs of both the CBM and the CBS, as described herein, but the embodiment is not so limited. The FIFO of one embodiment provides synchronous parameterized first-in first-out memory functions for bi-directional traffic. **Figure 29** is a  
20 parameterized bi-directional synchronous cell FIFO 2900 with loopback and processor access interface implemented in the CBC of one embodiment. A two-port REG\_FILE 2902 is used as the memory block, wherein the REG\_FILE 2902 is synchronous write and asynchronous read. An asynchronous interface is used for communication between the read 2904 and the write  
25 domain 2906, wherein the read ports and the write ports of each uni-directional FIFO are asynchronous to each other. The read and write logic is controlled by the clock of their respective domains 2904-2906. Programmable

parameters are used to customize the FIFO 2900, wherein the cell size, or number of words per cell, and the word size, or number of bits per word, are both programmable, or parameterized. Figure 30 lists the configuration parameters of the bi-directional FIFOs of the CBC of one embodiment. Figure 31 lists the input/outputs (I/Os) and a corresponding interface description for the bi-directional FIFOs of the CBC of one embodiment.

Specifically, the parameterized bi-directional FIFO unit controls cell traffic in a switch platform using a first and a second unidirectional FIFO buffer. The first and second unidirectional FIFO buffers each comprise asynchronous read and write ports. A cell size of the first and second unidirectional FIFO buffers is programmable. Furthermore, a word size of the first and second unidirectional FIFO buffers is programmable. The bi-directional FIFO unit is coupled to write at least one cell from and read at least one cell to at least one asynchronous transfer mode (ATM) interface, at least one frame relay interface, at least one voice interface, and at least one data interface, but the embodiment is not so limited. As such, the first unidirectional FIFO buffer is coupled to write at least one cell from an ATM interface, a frame relay interface, a voice interface, and a data interface. The first unidirectional FIFO buffer is coupled to read at least one cell to at least one switch, wherein the switch handles cells from sources having a number of bandwidths. The switch is coupled to route the at least one cell to an OC12 trunk line and to at least one service module. The service module is coupled to provide the cell to at least one service subscriber using T1, E1, T3, E3, OC3, and OC 12 ports, but the embodiment is not so limited.

The second unidirectional FIFO buffer is coupled to read at least one cell to an ATM interface, a frame relay interface, a voice interface, and a data interface. Furthermore, the second unidirectional FIFO buffer is coupled to

write at least one cell from at least one switch, wherein the switch handles cells from sources having a number of bandwidths. The switch may be coupled to route the cell from an OC12 trunk line and from a service module. The service module is coupled to provide the cell to at least one service subscriber using T1, E1, T3, E3, OC3, and OC 12 ports, but the embodiment is not so limited. Invalid cells may be discarded from each unidirectional FIFO buffer, but the embodiment is not so limited.

Figure 32 is a flowchart for controlling cell traffic in a switch platform. Operation begins at step 3202, at which a cell size of each of a first and second unidirectional FIFO buffers is programmed. A word size of each of the first and second unidirectional FIFO buffers is programmed, at step 3204. At least one cell is asynchronously transferred among a number of ports having a number of bandwidths, at step 3206. The asynchronous transfer uses a bidirectional first-in-first-out (FIFO) unit, wherein the bidirectional FIFO unit comprises a first and a second unidirectional FIFO buffer. Invalid cells are discarded from each unidirectional FIFO buffer, at step 3208. A non-destructive read is executed of the bidirectional FIFO unit, at step 3210, while cells are being written to and read from the bidirectional FIFO unit.

The step of asynchronously transferring may comprise the steps of synchronously writing the cell from a port to the first unidirectional FIFO buffer, and synchronously reading the cell from the first unidirectional FIFO buffer to at least one switch, wherein the reading is asynchronous with the writing, but the embodiment is not so limited. Furthermore, the step of asynchronously transferring may comprise the step of synchronously writing at least one cell from the switch to the second unidirectional FIFO buffer, and synchronously reading the cell from the second unidirectional FIFO buffer to



a port, wherein the reading is asynchronous with the writing, but the embodiment is not so limited.

In one embodiment, two signals pass through the asynchronous interface, a write\_port\_increment\_read\_cell\_counter signal, and a  
5 read\_port\_decrement\_write\_cell\_counter signal. As both of these signals are pulses, an asynchronous edge finder is used on each of these signals, but the embodiment is not so limited. The asynchronous edge finder of one  
10 embodiment is implemented in three steps. The first step is to make sure the pulse is long enough to be captured. If the signal is going from a fast-clocked register to a slow-clocked register, in order to ensure that the slow-clocked register captures the pulse, the fast-clocked register holds the value for a  
certain number of cycles. The number of cycles that the value is held is the ratio of the fast clock to the slow clock. For example, if the fast clock is 50MHz (20 ns period) and the slow clock is 21 MHz (48 ns), then the signal needs to be  
15 held for three cycles ( $20\text{ ns} * 3 = 60\text{ ns} > 48\text{ ns}$ ). This would ensure that the slow-clocked register captures the signal. Signals going from a slow-clocked register to a fast-clocked register do not have to deal with this issue. The second step handles meta-stability using doubling-latching logic. The third step is to capture the edge, wherein simple logic is used to make sure the  
20 signal is valid for one cycle in the destination clock domain.

The CBC of one embodiment supports high level flags comprising cell available flags and cell space available flags. **Figure 33** is a block diagram of the flag generation of one embodiment. The flags are based on cell-boundaries rather than word-boundaries, but the embodiment is not so  
25 limited.

A write port of the FIFO uses two counters: a write\_word\_counter, and a write\_cell\_counter. The write\_word\_counter counts the number of words

that have been written in a cell. The write\_cell\_counter counts the number of cells currently in the FIFO. When a whole cell is written into the FIFO, the write\_word\_counter resets to zero, the write\_cell\_counter is incremented by one. A signal is sent through the asynchronous interface to notify the read  
5 port to increment its read\_cell\_counter. The space\_available flag is based on cell boundaries. If there is one cell space left in the FIFO, as soon as the first word is written into the FIFO, the space\_available signal will be turned off. If the FIFO is full, the illegal\_write signal will prevent the write\_address register from incrementing.

10 In operation, each unidirectional FIFO buffer of one embodiment outputs a write port cell count from a write port. In one embodiment, a cell count value is programmed at which the write port cell count is outputted. When the write port cell count indicates that the FIFO buffer can not accept additional data or cells, the master bidirectional FIFO unit ceases reading cells  
15 to the unidirectional FIFO buffer of the slave bidirectional FIFO unit in response to the write port cell count. Furthermore, the master bidirectional FIFO unit disables a corresponding switch from routing cells to the slave bidirectional FIFO unit in response to the write port cell count; the switch routes the cells to another of the slave bidirectional FIFO units. The write  
20 port logic of each unidirectional FIFO buffer is synchronous with a write clock, wherein the write clock operates at a frequency substantially equal to 50 megahertz, but the embodiment is not so limited.

The master bidirectional FIFO unit resumes reading cells to the second unidirectional FIFO buffer of the slave bidirectional FIFO unit in response to  
25 a write port cell count indication that the FIFO buffer can accept additional data or cells. In response, the master bidirectional FIFO unit enables an associated switch to route cells to the slave bidirectional FIFO unit.

In addition to the write port cell count, each unidirectional FIFO buffer of one embodiment outputs a status indicating space available in the corresponding unidirectional FIFO buffer for at least one more cell. When the space available status indicates that space is not available in the FIFO buffer for at least one more cell, the master bidirectional FIFO unit ceases reading cells to the unidirectional FIFO buffer of the slave bidirectional FIFO unit in response to the space available status. Furthermore, the master bidirectional FIFO unit disables a corresponding switch from routing cells to the slave bidirectional FIFO unit in response to the space available status; the switch routes the cells to another of the slave bidirectional FIFO units.

The master bidirectional FIFO unit resumes reading cells to the second unidirectional FIFO buffer of the slave bidirectional FIFO unit in response to a status indicating space is available in the FIFO buffer. In response, the master bidirectional FIFO unit enables an associated switch to route cells to the slave bidirectional FIFO unit.

**Figure 34** is a timing diagram of a write port of one embodiment. The write\_enable 3402 and write\_data 3404 inputs are registered on the write clock 3406. An internal write\_address points to the next address to be written.

Each unidirectional FIFO buffer of one embodiment outputs a read port cell count from a read port. The read port cell count is a status indicating space available in the corresponding unidirectional FIFO buffer for at least one more cell. The read port logic of each unidirectional FIFO buffer is synchronous with a read clock. The read clock may operate at a frequency substantially equal to 21 megahertz or 42 megahertz, but the embodiment is not so limited.

A read port of the FIFO uses two counters, a read\_word\_counter, and a read\_cell\_counter. The read\_word\_counter counts the number of words

that have been read in a cell. The read\_cell\_counter counts the number of cells currently in the FIFO. When a whole cell is read from the FIFO, the read\_word\_counter resets to zero, and the read\_cell\_counter is decremented by one. A signal is sent through the asynchronous interface to notify the write port to decrement the write\_cell\_counter. The cell\_available flag is on cell boundaries. If there is one cell left in the FIFO, as soon as the first word is read from the FIFO, the cell\_available signal is turned off. If the FIFO is empty, the illegal\_read signal will prevent the read\_address register from incrementing.

**Figure 35** is a timing diagram of a read port of one embodiment. The read\_enable input 3502 is registered on the read clock 3504. The REG\_FILE already has the next read data available at its output. Once the registered read\_enable signal 3502 is active, it will load the data into the read\_data\_out register. This is the reason for the two clock delays between the read\_enable signal 3502 going active and data appearing on the read\_data\_out\_port.

**Figure 36** is a flowchart for providing cell bus management in a network switch of one embodiment. Operation begins at step 3602, at which a cell count value is programmed at which a first slave bidirectional FIFO unit outputs a write port cell count. A determination is made, at step 3604, as to whether the cell count of the slave bidirectional FIFO unit exceeds the programmed cell count value. If the cell count exceeds the programmed cell count value, the write port cell count is output from the first slave bidirectional FIFO unit. The reading of cells to the first slave bidirectional FIFO unit is suspended, at step 3606, in response to the write port cell count. Cells are routed to a second slave bidirectional FIFO unit in response to the write port cell count, at step 3608, in response to the suspended reading to the first bidirectional FIFO unit. The rerouted cell is read to the second slave

bidirectional FIFO unit in response to the write port cell count. If the write port cell count of the first slave bidirectional FIFO is determined not to exceed the programmed cell count value, or when the write port cell count drops back below the programmed cell count value, reading of the at least one cell to the first slave bidirectional FIFO unit is resumed, at step 3610. Cell bus management in the network switch of one embodiment further comprises outputting a read port cell count, and outputting from at least one read port a status indicating space available in a unidirectional FIFO buffer for at least one more cell.

Cell bus management in the network switch of one embodiment further comprises outputting a status indicating space available for at least one more cell from a first slave bidirectional FIFO unit when a cell count of at least one slave bidirectional FIFO unit exceeds one. Reading of at least one cell to the first slave bidirectional FIFO unit is suspended in response to a status indicating no space available. The cell is routed to a second slave bidirectional FIFO unit in response to a status indicating no space available. The rerouted cell is read to the second slave bidirectional FIFO unit. When the status of the first slave bidirectional FIFO unit again indicates space available in the buffer for at least one more cell, the routing of the at least one cell to the second slave bidirectional FIFO unit is disabled, and reading to the first slave bidirectional FIFO unit is resumed.

A processor port allows a processor to perform a non-destructive read for any address in the FIFO when such action is granted by the read port. A diagnostic interface of the bidirectional FIFO unit supports the non-destructive read of the bidirectional FIFO unit while at least one cell is being written to and read from the bidirectional FIFO unit. At least one cell may be written to the second unidirectional FIFO buffer from the first unidirectional

FIFO buffer over a first enabled diagnostic loop. Furthermore, at least one cell may be written to the first unidirectional FIFO buffer from the second unidirectional FIFO buffer over a second enabled diagnostic loop.

**Figure 37** is a timing diagram of a processor port of one embodiment.

- 5 The processor port is in the read port clock domain. Once the grant signal and the request signal are active, the processor port read address is substituted for the read port address. After the read is complete, the address register is switched back to the read port address.

10 In arbitrating slave module requests, the CBM performs arbitration on the requests from the eight slave modules based on a minimum guarantee bandwidth which is set by the ACP for each slave module. A particular slave module will get higher priority if it is not providing the corresponding minimum guarantee bandwidth.

15 In one embodiment, fair rate-based cell traffic arbitration is provided between eight traffic sources in the form of eight cell bus service modules on the same cell bus, but the embodiment is not so limited. A processor can program the CBC of one embodiment with an 8-bit Relative Service Delay (RSD) value for each of the eight service Modules. This RSD value determines the portion of the total bandwidth of the switch platform reserved  
20 for the respective service module. Given the bandwidth requirement for each of the eight service modules, the corresponding RSD values can be calculated using the formula

$$RSD_i = \frac{1/BW_i}{\sum_{n=1}^8 1/BW_n},$$

- 25 wherein  $RSD_i$  is the relative service delay for the  $i$ th service module, and  $BW_i$  is the bandwidth requirement for the  $i$ th service module. In one

embodiment, a bandwidth allocation for a service module is performed according to the formula

$$BW_i = \frac{1/RSD_i}{\sum_{n=1}^8 1/RSD_i},$$

5 wherein RSD<sub>i</sub> is the RSD value for the service module, but the embodiment is not so limited. If the bandwidth is under-subscribed, the remaining bandwidth is shared among all eight service modules according to the RSD value of each service module. If the bandwidth is over-subscribed, each service module will have the assigned bandwidth portion decreased  
10 according to the RSD values.

The rate-based cell traffic arbitration scheme of one embodiment is analogous to a parallel resistor network comprising eight resistors. **Figure 38** is a bandwidth distribution of one embodiment using a parallel resistor network model. The current 3802 flowing into the resistor network 3804 can  
15 be thought of as the total bandwidth supported by the CBM. The resistance value of each resistor 3810-3824 can be thought of as the RSD value for each service module. The current that flows across each resistor can be thought of as the portion of the total bandwidth allocated for the service module.

In operation, each service module comprises an 8-bit RSD register that  
20 is configured by the system processor, or ACP, using the RSD value. The value for each RSD is calculated based on the bandwidths allotted for each service module. Furthermore, each service module uses an 8-bit Service Delay Accumulator (SDA) register. The SDA register of each service module is configured using the SDA value, wherein the SDA register keeps track of  
25 when each of the service modules should receive service.

The SDA value is calculated at each cell bus frame time for each of the service modules based on the RSD value for each of the service modules, a request for service, and the minimum SDA value among the service modules during a cell bus frame time. In one embodiment, arbitration logic searches the SDA values of each of the service modules for a minimum SDA value, wherein the service module having the minimum SDA value and requesting service is selected for service. The SDA register of the service module having the minimum SDA value is configured with the RSD value when the service module is serviced. Each of the SDA registers of each of the service modules that do not have the minimum SDA value are configured with the result of the SDA value of the service module minus the minimum SDA value. In one embodiment, the SDA value is calculated for each clock cycle, but the embodiment is not so limited.

**Figure 39** is a flowchart for providing cell traffic arbitration in a switch platform of one embodiment. Operation begins at step 3902, at which at least one cell bus controller is programmed with at least one delay value for each of a number of service modules. The delay value determines a portion of a total bandwidth of the switch platform reserved for each of the service modules. The delay value comprises a relative service delay (RSD) value and a service delay accumulator (SDA) value, but the embodiment is not so limited. An 8-bit RSD register of each of the service modules is configured, at step 3904, using the RSD value. An 8-bit SDA register of each of the service modules is configured, at step 3906, using the SDA value, wherein the SDA register keeps track of when each of the plurality of service modules receives service. A determination is made, at step 3908, as to a portion of the total switch bandwidth to be reserved for each service module. When the total switch bandwidth is under-subscribed, the remaining bandwidth is shared



among the number of service modules. When the total switch bandwidth is over-subscribed, a portion of the total bandwidth allocated to at least one service module is decreased according to the RSD value.

The SDA value is calculated for each of the service modules at each cell bus frame time based on the RSD value for each of the service modules, a request for service, and the minimum SDA value among the service modules during a cell bus frame time. The SDA values of each of the service modules are then searched for a minimum SDA value, wherein the service module having the minimum SDA value and requesting service is selected for service. The SDA register of the service module having the minimum SDA value is configured with the RSD value when the service module is serviced. Each of the SDA registers of each of the service modules that do not have the minimum SDA value are configured with the result of the SDA value of the service module minus the minimum SDA value.

The CBC of one embodiment comprises diagnostic logic that assists in carrying out diagnostic functions. The diagnostic logic supported comprises: an interface loopback; a cell bus master-slave loopback, wherein a CBS loops back to any one CBM in the ingress direction, and the same selected CBM loops back to a CBS in the egress direction; an individual cell bus loopback; and an enable control for individual modules and functions.

**Figure 40** is a block diagram of an interface loopback of one embodiment. When the interface loopback mode is enabled, incoming cells will not be processed by the CBC but, instead, will be looped back from the ECIC 4002 to the ICOC 4004 and be sent back to the QE. Parity checking of incoming cells will be performed in the ECP 4003.

**Figure 41** is a block diagram of an individual cell bus loopback of one embodiment. When the individual cell bus loopback is selected, each cell bus

will loopback onto itself. Basically, the cell from the ECP 4102 will be looped back to the ICP 4104. When this loopback is selected, five cell buses go into loopback together.

Several functions of the CBC may be disabled in order to stop cell traffic through the CBC. Cell traffic is stopped so that the CBM and CBS FIFOs and queues can be monitored for cell progression through the CBC. The following modules and functions may be selectively disabled, but the embodiment is not so limited: the transmission of cells to the QEs by the ICOC; the transmission of cells to the ICOC by the MCE; the transmission of cell bus cells from the CBMs to the corresponding cell buses; and, the transmission of cell bus cells from the CBS to the slave cell bus.

While the present invention has been described with reference to specific exemplary embodiments, it will be evident that various modifications and changes may be made to these embodiments without departing from the broader spirit and scope of the invention as set forth in the claims. Accordingly, the specification and drawings are to be regarded in an illustrative rather than a restrictive sense.